

Particle swarming of sensor correction filters

Jonathan J Carter^{1,2,8}, Samuel J Cooper³, Edward Thrift³, Joseph Briggs⁴, Jim Warner^{5,6}, Michael P Ross⁷ and Conor M Mow-Lowry³

¹ Max-Planck-Institute for Gravitational Physics, Callinstr. 38, 30167 Hannover, Germany

² Institute for Gravitational Physics, Leibniz Universität Hannover, Callinstr. 38, 30167 Hannover, Germany

³ Institute for Gravitational Wave Astronomy, University of Birmingham, Birmingham, B15 2TT, United Kingdom

⁴ SUPA Institute for Gravitational Research, Department of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, United Kingdom

⁵ LIGO Hanford Observatory, Richland, Washington 99352, United States of America

⁶ LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, United States of America

⁷ Center for Experimental Nuclear Physics and Astrophysics, University of Washington, Seattle, Washington 98195, United States of America

E-mail: jonathan.carter@aei.mpg.de

Received 30 June 2020, revised 7 August 2020

Accepted for publication 27 August 2020

Published 21 September 2020



CrossMark

Abstract

Reducing the impact of seismic activity on the motion of suspended optics is essential for the operation of ground-based gravitational wave detectors. During periods of increased seismic activity, low-frequency ground translation and tilt cause the Advanced LIGO observatories to lose ‘lock’, reducing their duty cycles. This paper applies modern global-optimisation algorithms to aid in the design of the ‘sensor correction’ filter, used in the control of the active platforms. It is shown that a particle swarm algorithm that minimises a cost-function approximating the differential root mean squared velocity between platforms can produce control filters that perform better across most frequencies in the control bandwidth than those currently installed. These tests were conducted using training data from the LIGO Hanford Observatory seismic instruments

⁸ Author to whom any correspondence should be addressed.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

and simulations of the Horizontal Access Module Internal Seismic Isolation platforms. These results show that new methods of producing control filters are ready for use at LIGO. The filters were implemented at LIGO's Hanford Observatory, and use the resulting data to refine the cost function.

Keywords: sensor correction, seismic isolation, particle swarm

(Some figures may appear in colour only in the online journal)

1. Introduction

The first observation of gravitational waves from a binary black hole [1] was just a few years ago, but there has since been 11 confirmed astrophysical signals during the O1 and O2 observing periods [2], and many more event candidates during O3 [3] by the LIGO–VIRGO collaboration. For the Advanced LIGO detectors [4, 5], the increased frequency of detections has placed a greater emphasis on the stability and consistency of the observatories.

The extraordinary sensitivity of the detectors places stringent requirements on the residual vibration in the measurement band [6–8]. Additionally, for the observatories to be operational the suspended mirrors must be ‘locked’ to an ideal operating point by a feedback-control system [9–12]. If low frequency ground motion is insufficiently suppressed, it can cause the mirrors to move too far from the operating point. Losing and re-acquiring lock is a time-consuming process [8, 13, 14]. To reduce low-frequency vibrations, a complex multi-stage active isolation system is employed [15, 16]. The active control employs an array of differential and inertial sensors (ISs) on each stage. At frequencies between 0.3 and 15 Hz, feedback control provides most of the vibration suppression; near the microseismic peak (50–300 mHz) [17] it is largely achieved by feed-forward control.

‘Sensor correction’ (SC) is a control technique where signals from ground seismometers are filtered and summed with the output of differential displacement sensors that measure the position of LIGO's isolated platforms, as shown in figure 1. This removes the ground contribution of the displacement sensor to yield the inertial platform motion. The aim is to use the excellent noise seismometers able to detect motion even in a seismically quiet vault at low frequencies. However, there is a cut-off below 0.1 Hz where ground seismometers are dominated by their tilt susceptibility [18, 19]. This noise is avoided by using the displacement sensors as the feedback signal. A good SC filter aggressively tackles the translational ground motion without including frequencies where noise dominates, in particular, tilting contamination.

This paper presents a novel approach for improving sensor-correction filter design [20, 21] based on particle swarm optimisation. Improved filters are found with an unguided search using real data, for a variety of environmental conditions. By using real data-streams from all three types of installed sensors, and applying IIR filters exactly as LIGO's real-time control system does, the tool becomes immune to a slew of systematic issues including finite coherence, differing transfer-functions, and numerical implementation effects. A cost function was designed to mimic the physical quantity we believe is most important for LIGO's performance, the root mean squared (RMS) inertial velocity. The velocity spectrum was shaped in an *a priori* fashion to account for known physical effects of common-mode rejection and resonances of the suspended optics. Finally, the search variables were reparameterised to remove degeneracy in the search space.

In section 2, comparisons between current filter designs methods and this technique are discussed. Section 3 details the implementation of the particle swarm. Section 4 reviews the filters and their performance when installed at the Hanford site.

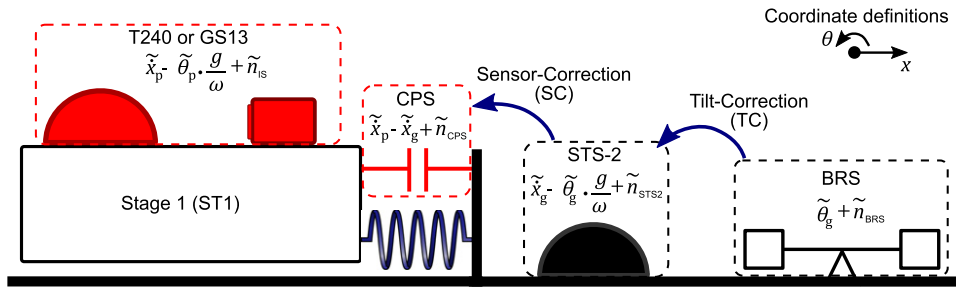


Figure 1. The signal path of tilt- and sensor-correction. Tilt correction is used at LLO in the beam-direction for isolation platforms in the corner station, while it is only used at LHO for the end stations (ETMX, ETMY). SC is used on all chambers at both observatories in translational degrees of freedom (X, Y, Z). The roles of each sensor are described in table 1. Figure adapted from Cooper [28] with permission.

2. Numerically optimised filter design

Filter design is a challenge in many fields; each of which have their own relevant filter design methods. The SC filters used at LIGO are hand-tuned by the operators and commissioners. While this can be an effective design method, it requires substantial experience and an understanding of the performance requirements. Until now there has not been a quantitative performance metric, and it has been difficult to evaluate the relative quality of different filters. Furthermore, there exists the possibility that the designs are limited by human biases. Additionally, since this design and testing process is time consuming, the filters are only replaced if there are obvious performance issues.

Other fields have developed a wide array of tools for generating of control filters to suit their specific needs. An example often used for motion suppression in car suspensions is reinforced learning automata, such as described by Howell *et al* [22] and other later groups [23]. Practical limitations prevent this method from direct application to gravitational wave detectors. It would require in operation testing so that the direct output of the filter can be evaluated. In operation measurements are not feasible for a gravitational-wave detector, due to the unacceptable loss of observation time of the testing process; this strongly motivates the development of rigorous off-line testing procedures.

Storn [24] pioneered a method for when there is *a priori* knowledge of the desired filter response. They use a differential evolution optimiser to fit poles and zeros to a prescribed phase and magnitude response, whilst employing cost ‘penalties’ in order to meet stringent stability conditions. This method is not well suited for designing a sensor-correction filter as the desired response is unknown due to the complex noise dynamics and limited coherence of the sensors in the isolation system [25].

Other groups have used particle swarms to design control filters based upon Storn’s work [26, 27]. They show that the particle swarm optimisation algorithm can effectively design filters due to its ability to sample large parameter spaces. A direct comparison between genetic algorithms and particle swarms is presented in [26]. The particle swarm reached a lower minimum more consistently, though there was little difference between the two methods; a similar conclusion was reached during the development of the optimisation method presented here. Modifications to the particle-swarm optimiser are described in [27] that made it similar to the ‘branch and prune’ optimisation technique. The changes significantly improved the

optimisation performance, and this illustrates the relative flexibility and adaptability of the particle-swarm technique.

A numerical optimiser enables a conceptually simple way to design SC filters. The only bound on the solver are the asymptotes outside the band of interest (4 mHz–4 Hz), which are entirely controlled by the number of poles and zeros in the filter. These asymptotes have strong physical motivation, however. Within the entire frequency band of interest a completely unguided and randomly seeded search is made. This allows for non-obvious, but beneficial features to be incorporated in the filter. Although conceptually simple, implementing such an optimiser proved to be a non trivial task.

Data collected from the wide array of on-site sensors was used for the calculation of filter performance during generation. Calculating the overall ground injection using this numerical, time domain approach allows for the inclusion of finite coherence between multiple sensors preventing the over-estimation of the SC filter. Furthermore, many real world effects of implementing filters such as differing transfer functions between sensors are already incorporated in the resultant filters. Because the data processed is in the time domain, only real coherence can be subtracted. Finally, doing so with time domain data is the numerically closest operation to on site action. This means an unstable filter will be inherently penalised as it will cause a large increase motion in the system.

While this work focuses on SC filters, the tools modular design means that it can be adapted to design various control filters within the seismic isolation system. The problem for filter generation then becomes purely a matter of finding an appropriate cost function.

3. Guiding the swarm

All numerical optimisation processes require a cost function. Using a particle swarm optimisation routine allows for a computationally intensive one. The cost function mimics the implementation of the sensor-correction filter performed by the LIGO's real-time control system. From there, the filter's performance was analysed to estimate the disturbances injected into LIGO's isolated platforms. Additional weighting was applied to account for key features such as suspension resonances.

3.1. Generating sensor-corrected signals off-line

Figure 1 shows a functional schematic of the SC signal path. Each block represents a type of sensor installed in the LIGO observatories and shows the contributions to their output signals in terms of inertial translation, tilt (or rotation), and self-noise. All variables carry a tilde to indicate the Fourier transform. From this block diagram, the sensor-corrected capacitive position sensor (CPS) signal, \tilde{a}_{sc} , can be calculated based on the signals from the seismometer and rotation sensor on the ground

$$\tilde{\theta}_r = \tilde{\theta}_g - \text{TC}(\tilde{\theta}_g + \tilde{n}_{\text{BRS}}) \quad (1)$$

$$\tilde{a}_{sc} = \tilde{\dot{x}}_p - \tilde{\dot{x}}_g + \tilde{n}_{\text{CPS}} + \text{SC}(\tilde{\dot{x}}_g - \tilde{\theta}_r \frac{g}{\omega}), \quad (2)$$

$$\tilde{a}_{sc} = \underbrace{\tilde{\dot{x}}_p}_{\text{platform motion}} + \tilde{n}_{\text{CPS}} - \underbrace{(1 - \text{SC})\tilde{\dot{x}}_g}_{\text{ground injection}} - \underbrace{\text{SC}(\tilde{\theta}_r \frac{g}{\omega} + \tilde{n}_{\text{STS2}})}_{\text{tilt injection}}. \quad (3)$$

Here \tilde{x} denotes translational motion, $\tilde{\theta}$ tilt motion, while subscripts g and p show whether this is associated with ground or platform motion respectively. The complex frequency responses of the sensor-correction and tilt-correction filters are denoted by SC and TC respectively,

Table 1. The instruments used in the construction of the SC cost. The names of the sensors are: BRS—beam rotation sensors, STS2—Streckeisen STS-2 force-feedback seismometer, CPS—capacitive position sensor, T240—trillium T240 force-feedback seismometer, and GS13—Geotech Instruments GS-13 short-period seismometer.

Sensor Name	Measures	Contributions
BRS	Ground inertial rotation	$\tilde{\theta}_g + \tilde{n}_{\text{BRS}}$
STS2	Ground inertial translation	$\tilde{x}_g + \tilde{n}_{\text{STS2}} + \tilde{\theta}_g \frac{g}{\omega}$
CPS	Relative platform velocity	$\dot{\tilde{x}}_p - \dot{\tilde{x}}_g + \tilde{n}_{\text{CPS}}$
T240	BSC-platform inertial translation	$\tilde{x}_p + \tilde{n}_{\text{T240}} + \tilde{\theta}_p \frac{g}{\omega}$
GS13	HAM-platform inertial translation	$\tilde{x}_p + \tilde{n}_{\text{GS13}} + \tilde{\theta}_p \frac{g}{\omega}$

ω is the angular frequency, and g is acceleration due to gravity. The ‘residual tilt’, $\tilde{\theta}_r$, is equal to the ground tilt in degrees of freedom where there is no BRS, as described below. The \tilde{n} terms denote the self-noise for the instrument indicated in the subscript. Table 1 describes the role of each instrument and shows the contributions that make up the output signals.

The ground and platform inertial motions are measured by seismometers. The ground rotation can only be measured where there is a beam rotation sensor (BRS) [29]. At the time of writing, there were four BRS instruments at LIGO Livingston Observatory (LLO), measuring the rotation that affects degrees of freedom parallel to the beam-line direction at the corner-station and at both end stations. At LIGO Hanford Observatory (LHO) there were two BRS instruments, one in the beam-direction at each end station.

In the ideal case, the sensor-corrected CPS signal only includes the platform motion, which can subsequently be suppressed by feedback control. By constructing the sensor-corrected CPS signal offline from 2 h stretches of data, many different sensor-correction filters can be implemented and tested without requiring valuable operating time of the LIGO observatories.

The optimisation problem of the sensor-correction filter is now clear: we must add as much of the inertial ground-motion term as possible to the CPS signal, while minimising the injection of tilt and inertial-sensor self-noise. The compromise between the two terms in equation (3) is complicated further by the small frequency separation of important dynamics in their spectra. The dominant contribution to the RMS ground translation occurs between 0.1 and 0.3 Hz (the secondary micro-seismic peak), and the tilt-injection begins to dominate the seismometer signals below approximately 0.05 Hz [30–32].

The SC filter can be thought of as one part of a complementary pair of ‘blending’ filters [33], an example of SC filters can be found in figure 6. The sensor-correction filter attenuates tilt and IS noise, while the ground injection is reduced by a factor of one minus the sensor-correction filter as can be seen from equation (3). When installed into the detector only the SC filter is installed, the filter complement is used to show the ground injection suppression.

The optimisation is restricted to a band between 4 mHz and 4 Hz. Below 4 mHz, the effect of all ISs is negligible on the platform RMS velocity [32]. Above 4 Hz, the CPS contribution to the platform feedback signal becomes negligible [15, 16], due to the suppression provided by the high and low pass blending filters respectively.

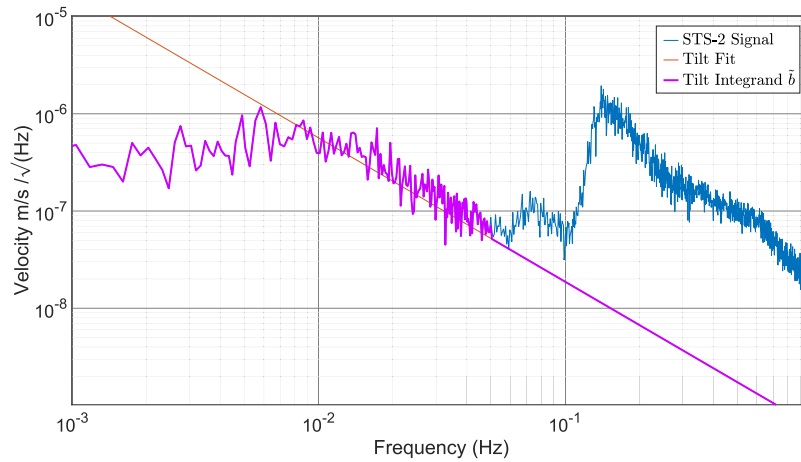


Figure 2. Figure showing the estimation of the tilt integrand, \tilde{b} , used in the swarming process. Between 0.01 Hz and 0.05 Hz the STS-2 motion is used, as this is strongly correlated with the local wind speed. A linear fit is used in this region to estimate the tilt coupling for the rest of the frequency band.

3.2. A physically motivated cost function

The design of a cost function often encodes much of the complexity of an optimisation problem. To aid this process, we created a cost that was the integral over frequency of the tilt-injection and ground-injection terms identified in equation (3). The ability to directly compare these spectral integrands with spectra of the input signals was of great practical benefit during debugging, and helped to shape both the cost function and performance penalties described below. The final cost is also approximately equal to the residual RMS velocity of the platform.

The RMS velocity of the platform is a useful figure of merit (FOM) because it is correlated with the scattered light performance of the interferometer [34]. It also balances the need to control drift at low frequencies (better evaluated by the RMS position) and limiting control forces at high-frequencies (better evaluated by RMS acceleration).

3.3. Tilt injection

At frequencies below approximately 50 mHz, the ground seismometer signal becomes completely dominated by tilt caused by wind pressure acting on the buildings [32]. Furthermore, the STS2 is AC-coupled at 8 mHz, and its output is no longer an inertial measure of the ground velocity. Therefore the STS2 signal below 50 mHz is considered to be part of the tilt-injection spectrum. Above this frequency, a simple power-law extrapolation of the tilt spectrum is made, as shown in figure 2. The small region from 50–100 mHz, which contains the primary micro-seismic peak, contains little spectral power and is de-weighted by the ground weighting function shown in figure 3.

The contribution to the total cost function is therefore the integral of the tilt-injection spectrum multiplied by the frequency response of the sensor-correction filter.




Figure 3. Weighting function, W_g , for the ground-injection spectral contribution to the cost function. Peaks at 0.45 Hz and 1 Hz correspond to suspension resonances.

3.4. Ground injection

With no sensor-correction filter applied, the inertial ground motion can be completely transferred to the platform through the CPS sensor. It is therefore important to reduce this term as much as possible, especially at the secondary micro-seismic peak, seen between 0.1 and 0.3 Hz in the plots below.

Each isolation platform has high Q , multi-stage suspensions that further reduce vibration above approximately 1 Hz. The lowest frequency pendulum-modes of the ‘quadruple’ suspensions, at 0.45 Hz and 1 Hz, are responsible for a significant fraction of the residual motion in LIGO, and it was important to include these dynamics in our cost. The weighting shown in figure 3 includes a simplified and broadened fit of the suspension response [35].

At frequencies below 0.1 Hz the ground injection cost is linearly de-weighted to zero. There are two reasons for this: first, the spectrum of the ground seismometer is increasingly dominated by tilt, and second, actual inertial translation is almost completely common-mode, even down the 4 km long arms of LIGO. A vanishingly small actuation force is needed to combat the differential component. At frequencies above 1 Hz, the CPS sensor plays very little role in the motion of the platform, and the cost is de-weighted to improve the numerical precision of the optimisation in the critical 0.1–1 Hz region.

There are therefore two components to the final cost of the swarming process, the ground and tilt injection. The ground injection signal is the previously calculated sensor corrected CPS signal minus the inertial platform motion, \tilde{x}_{inj} . This term is calculated in the time domain to account for the finite coherences between the ground seismometer, CPS, and IS on the platform. The tilt injection term is calculated in the frequency domain due to the estimation processed used to calculate the tilt integrand, \tilde{b} above 50 mHz. This is multiplied by the frequency response of the SC filter producing the tilt injection term.

The ground injection term, \tilde{x}_{inj} , is given by equation (4) and is defined here in the frequency domain. \tilde{a}_{sc} is the previously calculated sensor corrected CPS, \tilde{x}_p is the platform motion, \tilde{n}_{IS} is the noise of the IS, this is either a GS13 or a T240 and $\tilde{\theta}_p$ is the platform tilt. In reality, the swarming process calculates this in the time domain to ensure the swarmed filter is stable. The

tilt injection term, $\tilde{\theta}_{\text{inj}}$, is given by,

$$\tilde{x}_{\text{inj}} = \tilde{a}_{\text{sc}} - \dot{\tilde{x}}_{\text{p}} + \tilde{n}_{\text{IS}} + \tilde{\theta}_{\text{p}} \frac{g}{\omega}, \quad (4)$$

$$\tilde{\theta}_{\text{inj}} = \text{SC } \tilde{b}, \quad (5)$$

and is calculated in the frequency domain. Here the tilt fit curve, \tilde{b} , shown in figure 2, is multiplied by the SC filter and the ground injection term is multiplied by the ground weighting term, W_{g} , shown in figure 3. The tilt and ground injection terms are both integrated over the entire frequency band and summed in quadrature producing the final cost, given by

$$\text{cost}^2 = \int_{0.004 \text{ Hz}}^{4 \text{ Hz}} |W_{\text{g}} \tilde{x}_{\text{inj}}|^2 df + \int_{0.004 \text{ Hz}}^{4 \text{ Hz}} |\tilde{\theta}_{\text{inj}}|^2 df. \quad (6)$$

Using on-site data provides many advantages towards filter design. With the SC filter, the output of the platform seismometer acts as an *in situ* test of the quality of the filter. The goal of SC is to remove the ground motion in the CPS, such that the CPS follows the platform motion creating a quasi-IS. The T240 measures this and can therefore be used as a measure of quality, when factoring in the noise and tilt as measured by this sensor.

The designed filter can then be evaluated using other sets of input motion to verify their performance in a range of environmental conditions. Ensuring the cost function outlined above is physically motivated allows for quick comparisons between the performance of multiple filters and filters across different sets of environmental conditions.

3.5. Particle swarm optimisation

Particle swarm routines aim to simulate social behaviour to explore a parameter space. A number of ‘particles’ are generated with randomised parameters as part of the initial population. In the context of this work, a particle contain the information necessary to build a zero pole gain (ZPK) system that defines the sensor-correction filter. A ZPK function can be defined simply as a mathematical function of the following form

$$\frac{X_{\text{out}}}{X_{\text{in}}} = k \frac{\prod_{i=1}^I (s - z_i)}{\prod_{j=1}^J (s - p_j)}, \quad (7)$$

where X_{out} is the output, X_{in} is the input, z_i are the zeros, p_j are the poles and k is the system gain. The mapping from the degrees of freedom of the particles and the ZPK system is described in table 2. The particles are all then put through the cost function by the particle swarm routine. The cost function builds the ZPK of the filter from the particle’s parameters and then a cost for this particle is calculated using the relevant equations discussed in section 3.4. Particles then determine how to change their parameters for the next generation based on the overall global minimum cost, the local minimum, and a random ‘kick’. The process is then repeated until an exit condition of the swarm is met. During initial testing a swarm size of 500–1500 particles was able to sufficiently explore the parameter space whilst still enabling many iterations to be performed in a runtime less than 12 h on a typical laptop. Run time is linearly related to the number of particles as the run time is dominated by the number of amplitude spectral densities it must calculate for each particle.

To prevent overfitting to the training data the swarmed SC filter can only contain four complex poles and three complex zeros. The filters are seeded with a single real pole and three real zeros at 0 Hz to shape the filter and guarantee the required roll off at low frequency.

Table 2. A list of the parameters sent to the swarm and how they are converted into a filter function. Each frequency and Q is used to create one complementary pair of complex roots. The range of Q values allow the poles and zeroes to become slightly over-damped, making a pair of real poles.

Parameter	Description	Lower limit	Upper limit
1	Gain	0.95	1.05
2	$\log_{10}(f)$ of real pole	$\log_{10}(f_{\min})$	$\log_{10}(f_{\max})$
3	$\log_{10}(f_1)$ of highest complex pole	$\log_{10}(f_{\min})$	$\log_{10}(f_{\max})$
4–7	$\log_{10}(\Delta f_{1-3})$ from last complex pole	$\log_{10}(f_{\min}/f_{\max})$	0
8	$\log_{10}(f_1)$ of highest complex zero	$\log_{10}(f_{\min})$	$\log_{10}(f_{\max})$
9–11	$\log_{10}(\Delta f_{1-3})$ from last complex zero	$\log_{10}(f_{\min}/f_{\max})$	0
12–20	Quality factors for complex roots	0.35	5

The use of a particle swarm allows for an unguided search of the created parameter space. It is forced to construct the initial generation of filters with random parameters for the frequency and Q (for complex poles and zeros) so that the parameter space for any potential filters is fully explored.

3.6. Reparameterisation

How the parameter space is explored by the optimisation routine was the first area considered when refining the filter generation tool. By choosing the space in which the roots are generated, the weighting of filter construction in early generations can be skewed to allow for faster and more reliable convergence. The positions of roots in an ‘ideal’ filter will typically not be distributed uniformly. If the filters were generated in a linear frequency space, where the value generated is the position of the root, then shaping the low frequency response in a wide data range would take a much greater optimisation time. Instead, to account for various physical factors in certain frequency bands, roots will typically be spread over a logarithmic scale with a high density at crossover points to allow for better shaping.

The degeneracy of the space is the largest problem with a natural frequency creation, (where the value of the particle is the frequency of the root). Any two roots of a polynomial can be interchanged and the resulting transfer function will remain the same, displayed by figure 4 (top). This artificially inflates the parameter space by a factor proportional to the number of roots factorial. Larger parameter spaces will take longer for the optimisation routine to fully explore. To solve this problem, roots were ordered as shown in figure 4 (bottom). Here the lowest frequency root is defined in natural frequency space and the subsequent roots defined in terms of the difference to the last one generated. This orders the roots, solving the degeneracy problem. Ordering of the roots presents different problems to the filter construction in the form of imposing boundaries on root placement. To maintain the appropriate slopes outside of the studied region, no poles or zeros can be placed there. When building a filter, the range of acceptable poles and zeros should therefore be bound to the range of frequencies that the data covers. However, by defining the relations between roots, no boundary conditions can be imposed as shown by figure 4, where each jump is less than the range of frequencies allowed, but still the final root lies outside the acceptable range. This wastes a large amount of the available computing power on filters which must be discarded.

With effort on two fronts this problem was overcome. The first thing was to add a large scaling cost when an invalid filter was produced. This extra cost was based on a multiplicative term of how many roots and how far each root was outside the acceptable range. This created a ‘bucket’ in the parameter space which allowed particles to drift back into the acceptable range




Figure 4. (Top) the problem of roots degeneracy where f_1 and f_2 can be switched and the resultant filter remain the same. (Bottom) how the roots of a polynomial can escape the frequency range of desired generation with jumps of size less than the frequency range when ordered generation is used.

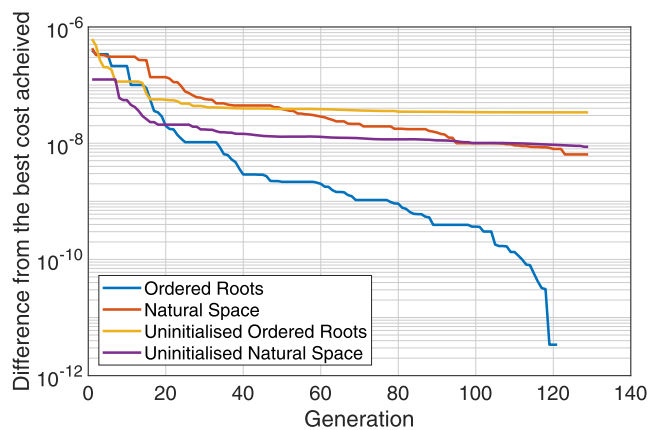


Figure 5. The convergence to best filter of the different parameter spaces (ordered and natural) and the effect of swarm initialisation. All data is measured relative to the best filter. This corresponds to a 13% improvement in frequency space. The tests were run several times and those with similar cost in initial generation displayed.

by imposing a gradient towards acceptability. The second tactic was a careful initialisation of the swarm matrix. Before the swarm was run, a first generation of filters was created in natural frequency space and then ordered from highest to lowest. The difference between these was then calculated and used as the parameters for generation in the swarm. This meant initial generation of particles were all viable filters and any particle which drifted out of range was pushed back by the aforementioned bucket. A graph showing the convergence of the different methods discussed is shown in figure 5, showing that after 120 generations with 500 particles the ordered generation initialised jumps achieved a significant improvement in best cost. It showed that both steps were necessary to see an improvement.

The parameters used by the swarm to make a control filter are shown in table 2.

4. Implementation at LIGO Hanford Observatory

4.1. Choice of training data

The data used for filter generation was collected during the commissioning break between the second and third observing runs, when the detector was in a damped only state (the isolation forward control loops were not in operation). This simplifies the sample data as the many control techniques did not have to be considered in the processing of the input data, nor was it necessary to account for closed loop stability requirements.

The swarmed filters were created using data collected during times of above average environmental disturbance (500 nm s⁻¹ bandwidth-limited RMS-velocity between 0.1 Hz and 0.3 Hz motion, and 8–9 m s⁻¹ windspeed). Filters trained and built during different environmental conditions are discussed in chapter 6 of Cooper’s thesis [28]. Since the detectors have the ability to switch between different filters, it would be possible to generate filters to account for seasonal variations of ground motion and wind speed, though this is outside the scope of the paper.

4.2. Optimising multiple chambers

In the corner stations of both LIGO observatories one sensor-correction filter is applied to the STS-2 signal in each degree of freedom, and distributed to all chambers. This ensures that the residual tilt injection that couples through the SC path is the same for each of the isolation platforms, reducing differential platform translational motion [28] (chapter 6). Therefore when optimising the sensor-correction filters for the corner station, a filter must be produced that is effective in all chambers. To produce the required input data for the swarming process, signals from each chamber on the ‘beam axis’ were averaged together. For each of the horizontal translational degrees of freedom (X and Y) we combine the signals from all the relevant chambers in the main building.

4.3. Comparison with current filter

The filters produced by the swarm were tested at the Hanford site between the 11th of October 2019 and 29th October 2019. The observatory had several other events occurring that resulted in the Internal Seismic Isolation (ISI) platforms operating at atmospheric pressure leading to overall higher noise. On site operators noted that the filter injected more ground tilt into the ISI, this is due to the filter having a factor 6 higher gain at low frequency compared to the previous SC filter.

The direct comparison of two SC filters that were active during two different stretches of time is not trivial due to the differences in the input ground and wind conditions. As a result, the cost function that was used to design the filters would not be appropriate. Moreover, the GS-13, used on the Horizontal Access Module (HAM) ISIs as a witness sensor is limited by its own self noise below approximately 0.05 Hz, is susceptible to platform tilt, and thus would not be a good witness sensor below this frequencies. To make a fair comparison, filter performance was evaluated during periods with similar environmental conditions. Both times had an average level of microseismic motion and low wind speed. To compensate for the remaining differences in input ground motion, the FOM term is weighted by the ratio of the amplitude spectral densities of the ground seismometer signals from the two testing periods.




Figure 6. Figure showing a comparison between the swarmed SC filter and current SC filter shown in red and blue respectively. The filter itself is shown by the dashed line and its complement is shown by a solid line. Reproduced with permission from Cooper [28].

As such, below approximately 0.1 Hz the performance of the filters was determined by the sensor corrected CPS signal and by the GS-13 or T-240 above this frequency. The exact form of the FOM is given by equation (9). The sensor corrected CPS provides information on the relative tilt injection and partially the suppression of the microseismic peak while the on-platform IS measures the suppression of the microseismic peak and the isolation performance above 0.1 Hz. It is important to quantify the FOM in terms of RMS velocity, as this is strongly correlated with scattered light in the interferometer moreover the velocity is a compromise between the need to control drift at low frequency (best evaluated by the RMS position) and limiting control forces at high-frequencies (best evaluated by the RMS acceleration).

To quantify the performance of a SC filter, the amplitude spectral density of the signals from each sensor described in table 1 was recorded. The spectra were standardised by converting them into velocity and normalising them to account for the frequency-dependence of their sensitivity, as appropriate. The spectra are then combined such that the CPS is used between 1 mHz and 0.1 Hz and the on-platform IS, is used between 0.1 Hz and 1.5 Hz. The IS is either a GS13 for the HAM chambers or a T240 for the BSC chambers. The resulting spectrum is defined as

$$\tilde{x}_m = \begin{cases} \text{CPS}, & 0.001 \leq f \leq 0.1 \\ \text{IS}, & 0.1 < f \leq 1.5. \end{cases} \quad (8)$$

This ‘combined’ signal is then integrated over the band as shown producing the FOM,

$$\text{FOM} = \int_{0.001}^{1.5} |\tilde{x}_m|^2 df. \quad (9)$$

A comparison of the swarmed SC filter and the current filter is shown in figure 6. The swarmed filter has significantly less gain peaking around 50 mHz resulting in more tilt injection below 10 mHz. Figure 7 shows a comparison between the two SC filters when evaluated on HAM2 in the X degree of freedom. The RMS of each of the traces are calculated from 1.5 Hz to prevent the large spikes at 1.7 and 2.2 Hz due to the suspension resonances from saturating the




Figure 7. The cost given by equation (9) for two different SC filters while running on HAM2 in the X degree of freedom. The RMS velocity of each of the filters is given by the dashed traces.




Figure 8. Figure shows the FOM given by equation (9) for two different SC filters while running on ITMX in the X degree of freedom. The RMS of each of the filters is given by the dashed traces.

RMS. Overall, the swarmed SC filter results in 27% lower RMS velocity of the ISI compared to the current filter when the input ground motion has been normalised. While the swarmed filter causes a factor 10 more tilt injection at 1 mHz this does little to affect the overall RMS of the platform. The swarmed filter makes most of its improvement between 0.4 and 0.9 Hz, of particular interest is the performance difference at 0.45 Hz, one of the suspension resonances, where the swarmed filter results in 33% less motion than the current filter.

Figure 8 shows a comparison between the swarmed and current SC filter on ITMX in the X degree of freedom. The swarmed filter results in a more modest 18% reduction in RMS velocity of the ISI. Again, the swarmed filter causes more tilt injection between 30 and

200 mHz, though this has little effect on the overall velocity RMS. The biggest difference between the two filters comes between 0.4 and 0.9 Hz. At 0.45 Hz the swarmed filter results in 68% less motion than the current SC filter.

5. Conclusions

Many strategies are being developed to decrease lock loss due to ground motions at the LIGO facilities. Improving the control filter design presents an opportunity to make significant gains with simple and affordable changes. A robust working tool to design SC filters has been developed, and filters created by this method have been deployed on-site. This work focused on one specific filter although its design principles can be adapted for other filters. The blend filters present in all the isolation systems are excellent candidates for testing. Designing control filters in this manner presented some challenges that had to be overcome, but techniques were developed, trailed, and tested throughout have solved many of the issues and may prove relevant to other applications. With this tool, better seismic isolation can be achieved leading to more observing time, more source detections and, ultimately, a greater understanding of some of the most fascinating phenomena in the Universe.

Acknowledgments

We would like to thank Brett Shapiro and Brian Lantz for their many useful suggestions and comments. Furthermore, discussions with Will Farr gave an insightful approach to building filters.

ORCID iDs

Jonathan J Carter  <https://orcid.org/0000-0001-8845-0900>
Samuel J Cooper  <https://orcid.org/0000-0001-8114-3596>
Michael P Ross  <https://orcid.org/0000-0002-8955-5269>
Conor M Mow-Lowry  <https://orcid.org/0000-0002-0351-4555>

References

- [1] LIGO Scientific Collaboration and Virgo Collaboration 2016 *Phys. Rev. Lett.* **116** 061102
- [2] LIGO Scientific Collaboration and Virgo Collaboration 2019 *Phys. Rev. X* **9** 031040
- [3] Cooper S, Morrell S, Jones A, Freise A and Smetana G 2019 Chirp gravitational wave alerts <https://zenodo.org/record/3525064>
- [4] LIGO Scientific Collaboration 2015 *Class. Quantum Grav.* **32** 074001
- [5] LIGO Scientific Collaboration and Virgo Collaboration 2016 *Phys. Rev. Lett.* **116** 131103
- [6] Robertson N A *et al* 2002 *Class. Quantum Grav.* **19** 4043–58
- [7] Barsotti L, Evans M and Fritschel P 2010 *Class. Quantum Grav.* **27** 084026
- [8] Staley A *et al* 2014 *Class. Quantum Grav.* **31** 245010
- [9] Martynov D V 2015 Lock acquisition and sensitivity analysis of advanced LIGO interferometers *PhD Thesis* California Institute of Technology <https://thesis.library.caltech.edu/8899/>
- [10] Datrier L and Dwyer S 2019 LIGO document control center portal <https://dcc.ligo.org/LIGO-T1900433>
- [11] Ross M 2019 Updated lock vs wind <https://alog.ligo-la.caltech.edu/SEI/index.php?callRep=1544>
- [12] Evans M *et al* 2002 *Opt. Lett.* **27** 598–600
- [13] Whitcomb A J, Slagmolen B J J, Miller J, Evans M, Fritschel P, Sigg D, Waldman S J, Shaddock D A and McClelland D E 2012 *Opt. Express* **20** 81–9

- [14] Acernese F et al 2006 *Class. Quantum Grav.* **23** S85
- [15] Matichard F et al 2015 *Precis. Eng.* **40** 287–97
- [16] Matichard F et al 2015 *Precis. Eng.* **40** 273–86
- [17] Longuet-Higgins M S and Jeffreys H 1950 *Phil. Trans. R. Soc. A* **243** 1–35
- [18] Lantz B, Schofield R, O'Reilly B, Clark D E and DeBra D 2009 *Bull. Seismol. Soc. Am.* **99** 980–9
- [19] Arai K (TAMA Collaboration) 2002 *Class. Quantum Grav.* **19** 1843
- [20] DeRosa R, Driggers J C, Atkinson D, Miao H, Frolov V, Landry M, Giaime J A and Adhikari R X 2012 *Class. Quantum Grav.* **29** 215008
- [21] Hua W 2005 Low frequency vibration isolation and alignment system for advanced LIGO *PhD Thesis* Stanford University https://labcit.ligo.caltech.edu/~rana/docs/Theses/hua_thesis.pdf
- [22] Howell M N, Frost G P, Gordon T J and Wu Q H 1997 *Mechatronics* **7** 263–76
- [23] Kashki M, Abido M A and Abdel-Magid Y L 2010 *Electr. Eng.* **91** 383–94
- [24] Storn R 1996 Differential evolution design of an IIR-filter *Proc. IEEE Int. Conf. on Evolutionary Computation* pp 268–73
- [25] Matichard F et al 2015 *Class. Quantum Grav.* **32** 185003
- [26] Krusienski D J and Jenkins W K 2004 Particle swarm optimization for adaptive IIR filter structures *Proc. 2004 Congress on Evolutionary Computation (IEEE Cat. No.04TH8753)* vol 1 pp 965–70
- [27] Luitel B and Venayagamoorthy G K 2008 Particle swarm optimization with quantum infusion for the design of digital filters *2008 IEEE Swarm Intelligence Symp.* pp 1–8
- [28] Cooper S J 2019 Breaking the seismic wall: how to improve gravitational wave detectors at low frequency *PhD Thesis* University of Birmingham <http://etheses.bham.ac.uk/id/eprint/9903>
- [29] Venkateswara K et al 2017 *Bull. Seismol. Soc. Am.* **107** 709–17
- [30] Schofield K 2016 Livingston ground motion under high winds <https://alog.ligo-la.caltech.edu/SEI/index.php?callRep=1089>
- [31] Venkateswara K, Hagedorn C A, Turner M D, Arp T and Gundlach J H 2014 *Rev. Sci. Instrum.* **85** 015005
- [32] Venkateswara K et al 2020 arXiv:2003.06447
- [33] Kissel J S 2010 Calibrating and improving the sensitivity of the LIGO detectors <https://dcc.ligo.org/LIGO-P1000103/public>
- [34] Ottaway D J, Fritschel P and Waldman S J 2012 *Opt. Express* **20** 8329–36
- [35] Aston S M et al 2012 *Class. Quantum Grav.* **29** 235004